JPL Document D-60951

Exoplanet Exploration Coronagraph Technology

Technology Milestone #2 Report

August 8, 2008

Editors:

Brian Kern, Andreas Kuhnert and John Trauger

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Jet Propulsion Laboratory
California Institute of Technology

EXOPLANET EXPLORATION CORONAGRAPH TECHNOLOGY

TECHNOLOGY MILESTONE #2 REPORT
BROADBAND STARLIGHT SUPPRESSION DEMONTRATION

Prepared by

Brian Kern, HCIT Experiment Lead, JPL — 9-11-08 Date

Approved by

Marie Levine, EXEP Coronagraph Technology Manager, JPL — 9/11/08 Date

Michael Devirian, EXEP Manager, JPL — 9/3/08 Date

Zlatan Tsetanov, EXEP Program Scientist, NASA HQ — 9/18/2008 Date

Lia LaPiana, EXEP Program Executive, NASA HQ — 9/18/08 Date

TABLE OF CONTENTS

1 ABSTRACT

This report discusses the achievement of TPF-C Technology Milestone 2 (M2), the Broadband Starlight Suppression Demonstration in May 2008. This is a laboratory demonstration of contrast better than 10^{-9}, observed through a 9% fractional bandwidth, over a region of the image plane relevant to planet-finding missions.

This report reviews the definition of the milestone from the Technology Milestone #2 White Paper (2006 Nov 28), reports the laboratory data, and addresses the statistical requirements of the milestone.

2 INTRODUCTION

A space mission intending to directly image exoplanets using a coronagraph will require significant new technologies, and the purpose of the High Contrast Imaging Testbed (HCIT) is to demonstrate a collection of these technologies at relevant levels of performance. In 2006, TPF-C Milestone #1 (M1) was met, which was specified in the M1 White Paper as follows: "*Demonstrate that the High Contrast Imaging Testbed (HCIT) achieves a baseline contrast of 1×10^{-9} (goal 1×10^{-10}) at a 4λ/D inner working angle, at a wavelength of 785 nm and stable for at least one hour.*"

Milestone 2 has a similar goal to Milestone 1, but requires a broader bandpass, as specified in the M2 White Paper: "*Demonstrate that the HCIT is capable of achieving a baseline contrast of 1×10^{-9} (goal 1×10^{-10}) at a 4 λ_0 / D inner working angle with a 10% spectral passband centered at a wavelength λ_0 in the range from 500 to 800 nm.*"

The rationale behind the M2 requirements is to demonstrate performance under particular conditions that represent challenges to a space mission. The inner working angle is a strong driver of telescope size, and therefore mission cost. The wavefront control of light at small angles is more difficult than for light at large angles, so the milestone requirement that this performance be demonstrated at small working angles tests the wavefront control in the regime that directly influences the feasibility and cost of a space mission.

The level of contrast specified in the M2 requirements is relevant to the goal of finding terrestrial planets around nearby stars. While photon statistics and exozodiacal light would be expected to limit detectability in a space mission, the speckle background must be kept low because the speckled nature of that background masks target planets. The variations among speckles will generally be some fraction of the mean background level, so requiring that the mean background level be comparable to the target brightness has relevance as a rough detectability threshold.

Large bandwidths are desirable for two main reasons: to minimize exposure times during imaging and to maximize spectroscopic coverage. A 10% fractional bandwidth is comparable to a typical color-filter bandwidth, and demonstration at any significant fractional bandwidth is likely to identify potential systematic difficulties associated with chromaticity.

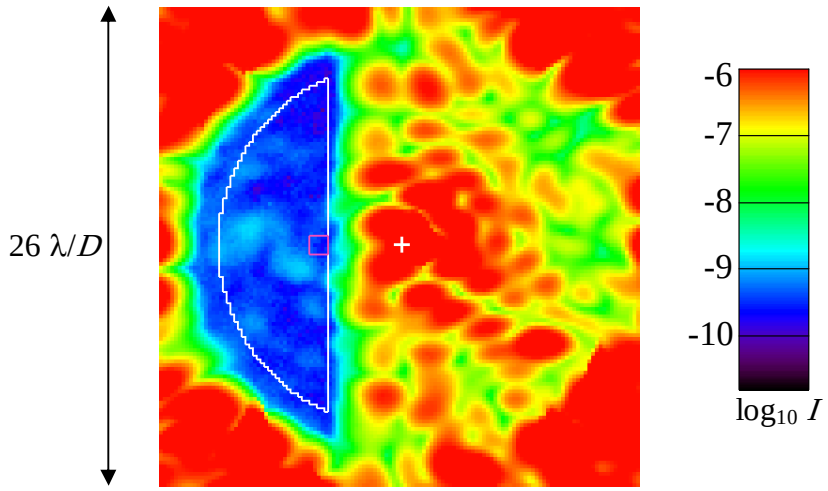

Fig. 1. Average coronagraph image, with 9% bandpass. This is the average of all iterations in "Set 3," defined in Section 5. The average contrast is 6.1×10^{-10} in the white "D" region and 4.7×10^{-10} in the small red box. The central star is marked with a "+". These contrast levels are well below the 10^{-9} milestone requirement. The field of view is ±13 λ/D.

As will be described in Section 5, the M2 demonstration contrasts exceeded the requirements by a significant margin. A sample image is shown in Fig. 1, which represents a typical dataset from this experiment (an average over many iterations). The average contrast, required to be below 10^{-9} in two image plane regions, is 6.1×10^{-10} and 4.7×10^{-10} in this sample. The bandwidth in which the contrast was measured was 9%, which is smaller than the 10% goal of M2. We have enough spectral information to extrapolate reliably what the contrast would have been in a 10% bandpass, and we show in Section 5.4 that a careful interpretation of the spectral data clearly implies that both goals are met in a 10% bandpass by a similarly large margin. We perform a phase-only correction in the pupil plane, which can correct phase and amplitude aberrations over only half of the image plane, to first order, which is why the M2 demonstration is performed over half the image. Future experiments with two deformable mirrors, one not appearing in the pupil plane, will be able to correct both sides of the image plane.

Much of the equipment and procedures used for the M2 demonstration is similar to what was used for the M1 demonstration. The significant changes are the broadband light source (Section 3.2), the optical properties of the occulting mask (Section 3.3), the wavefront control algorithm (Section 3.4), and the image normalization process (Section 4.2).

3 HCIT OVERVIEW

The milestone demonstration has been carried out with the High Contrast Imaging Testbed (HCIT), a laboratory facility presently located in the JPL Optical Interferometry Demonstration Laboratory. The testbed supports the development of two fundamental elements of the baseline exoplanet high-contrast imaging strategy. These are (1) suppression of scattered light via wavefront control with a single deformable mirror (DM) using images at the science camera, and (2) suppression of diffracted light via a Lyot coronagraph with a band-limited focal plane occulter. Other configurations are planned for future demonstrations, including wavefront control

with multiple DMs in the optical path and diffraction control using alternate coronagraph configurations.

HCIT represents ten years of JPL development in actively-controlled space coronagraph concepts, for the exploration of nearby exoplanetary systems. As a multi-purpose, rapid prototype development facility for coronagraph technologies, algorithms, and predictive models, HCIT uses as much as possible inexpensive off-the-shelf optical and electronic components. Initial experiments in 1998 involved prototypes of the modular DMs, based on lead magnesium niobate (PMN) electrostrictive technology at Xinetics and supporting control electronics and algorithms. These were mounted on a vacuum Michelson interferometer "surface gauge" capable of directly imaging the DM surface with 0.1 × 0.1 mm resolution (100 pixels per actuator) and 10 pm accuracy for surface deviations. This DM development, supported over the past ten years by a number of Small Business Innovation Research (SBIR) grants to Xinetics and technology funding at JPL, leveraged the PMN electrostrictive technology widely in use in discrete-actuator deformable mirrors at ground based observatories. Surface figure control demonstrations at the 25 picometer root mean square (rms) level, together with optical Fresnel propagation models for a variety of Lyot coronagraphs and speckle nulling algorithms subsequently indicated that the wavefront control technology was ready for demonstrations of high-contrast imaging in a dedicated coronagraph testbed (Trauger et al. 2002). Coronagraph elements were first assembled on the optical table in 2002. Initial experiments were carried out without a vacuum chamber in the ambient laboratory environment. The optical table (as well as the original Michelson "surface gauge") was moved to its present location and installed in the vacuum chamber in 2003. HCIT now facilitates high contrast demonstrations including outside experimenters from the larger exoplanet community outside of JPL. Upgrades of the testbed optics are anticipated for future high-contrast imaging demonstrations approaching the level of flight requirements.

3.1 HCIT Optical Layout

The coronagraph optics are mounted on a 5 ft × 7 ft optical table, which is supported on elastomeric vibration isolation blocks within a cylindrical vacuum chamber measuring 6 ft in diameter and 8 ft in length. The vacuum chamber is supported and vibration isolated from the laboratory floor by six pneumatic isolation legs, and during operation is evacuated to approximately 1 milliTorr with a dry scroll pump. The chamber walls are wrapped with a thermostatically controlled heating tape and a thermal blanket to maintain the chamber walls at temperature a few degrees above ambient. The optical layout is illustrated in Fig. 2.

With reference to Fig. 2, the elements located at upper right prior to the occulting mask perform the wavefront correction, relaying the pinhole light source to the occulting mask and ultimately correcting a $\lambda/10$ rms wavefront to $\sim \lambda/10^6$ rms over some spatial frequencies. The light source for this experiment is a supercontinuum laser source, described in more detail in Section 3.2.

The source "star" is reimaged on the coronagraph occulter with the following optical elements. OAP1 and OAP2 are commercial off-the-shelf mirrors from SORL with identical specifications (OAP 30-05.75-03.5SQ, 30 inch parent parabola focal length, 3.5 inch diameter, $\lambda/16$ peak-to-valley surface accuracy over 95% of the mirror). The fold flat has $\lambda/100$ rms surface quality. The

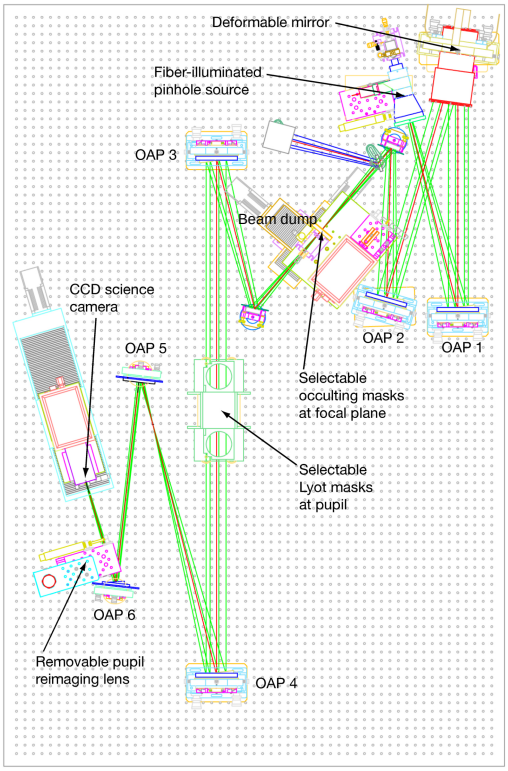

Fig. 2. Mechanical layout of HCIT Lyot coronagraph. Light enters the coronagraph at the "Fiber-illuminated pinhole source" near the top-right corner, reflects off the off-axis parabolic mirrors (OAPs) in numerical order (1–6), and lands on the "CCD science camera" on the left side of the figure. The deformable mirror (DM) enables wavefront control, and the occulting mask and Lyot mask create a Lyot coronagraph. All mirror surfaces except the DM have bare gold coatings. The occulting masks, Lyot masks, and science camera are mounted on motorized stages for remote selection and alignment.

deformable mirror is manufactured by Xinetics, Inc. It has a monolithic array of 1024 PMN electroceramic actuators, in a 32 × 32 square array, with an actuator pitch of 1 mm. The mirror facesheet is polished to λ/100 rms. A 30 mm diameter black-anodized aperture is placed in close proximity to the DM, defining the system pupil. The actuator driver system includes a control computer, a power supply, multiplexer driver, digital-to-analog converters (DACs) and clocking circuits which feed a set of low-power low-leakage multiplexer switch arrays in close proximity to the DM inside the vacuum chamber. Control computers, driver electronics, and thermal control hardware reside outside the vacuum chamber. The reflective surface on all mirrors is unprotected gold, with reflectances uniform over the surface at the level of a few percent.

The elements from the coronagraph occulter to the science camera perform the suppression of diffracted light with a Lyot coronagraph, described in more detail in Section 3.3. The beam is collimated by OAP3, which also projects an image of the DM and its pupil stop to the plane of the Lyot stop. OAP4 brings the light to an intermediate focus before OAP5. OAP3 and OAP4 are identical to OAP1 and OAP2. The final two off-axis parabolic mirrors OAP5 and OAP6 relay the image to the focal plane with a 3:1 magnification. The science focal plane is an e2V 1033 × 1056 pixel, back-illuminated and antireflection-coated CCD, cooled thermoelectrically to −70° C. All imaging in HCIT is performed at this single focal plane, including all wavefront sensing and correction and the recording of high contrast images. The only shutter in the system is located in air as part of the supercontinuum source; there is no shutter in the vacuum chamber.

3.2 Supercontinuum light source

The broadband light source used for the nulling experiments is a supercontinuum source built at JPL. Supercontinuum generation is the formation of broad, continuous spectra by propagation of

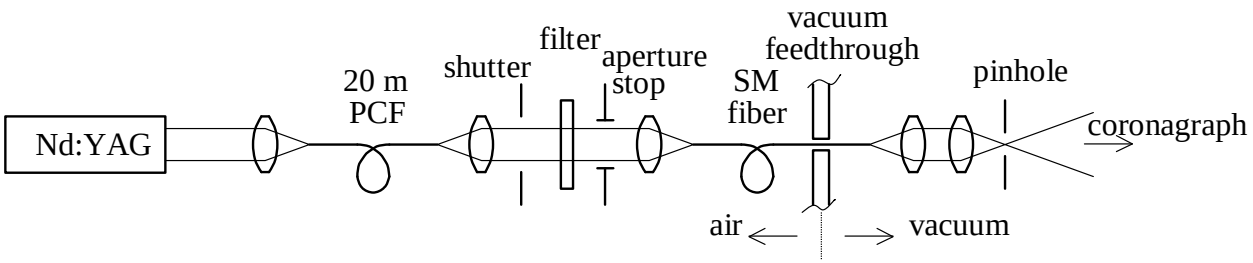

Fig. 3. Schematic supercontinuum source description. The output of a Nd:YAG laser is coupled to a 20 m long photonic-crystal fiber (PCF). The output of the PCF is collimated and shuttered, sent through a bandpass filter, stopped down, then coupled into a single-mode fiber (SM fiber). This fiber enters the vacuum chamber and is reimaged onto a 5 μm pinhole (a spatial filter). The output from the pinhole illuminates the coronagraph.

high power pulses through non-linear media (Alfano & Shapiro 1970a and b). The non-linear medium used in our case is a single-mode photonic-crystal fiber (PCF), model SC–5.0–1040 from Crystal Fibre, 20 m long. This is an index-guided PCF, comprising a high-index, solid glass core surrounded by a lower index air-filled cladding structure, *i.e.*, air holes arranged in a pattern that runs along the fiber. This PCF is pumped by a 1064 nm wavelength Nd:YAG passively Q-switched micro-chip laser, producing sub-nanosecond pulses with 6 μJ/pulse at a repetition rate of 7 kHz. This arrangement is shown on the left side of Fig. 3. The PCF output has a broad spectrum, with appreciable spectral power distributed from 600 to 1500 nm, with little net power loss from the pump laser, laterally confined in the fundamental mode of the fiber.

The broadband output of the PCF is collimated and sent through one of several bandpass filters. A total of six bandpass filters were available, on a motorized filter wheel, of which four were used for this experiment. The bandpass in which Milestone 2 is defined is a 10% fractional bandpass centered at 800 nm, which was measured through a bandpass filter with a transmission full width at half-maximum (FWHM) of 75.3 nm, mean wavelength 800.4 nm, for a 9.4% fractional transmission bandpass. The wavefront control algorithm makes use of additional bandpasses, each with a nominal transmission bandpass of 2%, centered at 768, 800, and 832 nm (*i.e.*, −4%, 0, +4% from 800 nm). We also image through two extra 2% bandpass filters at 784 and 816 nm (-2%, +2%), which are not used for scoring or for wavefront control but do provide contiguous spectral information over the 10% bandpass for additional investigations. Measurements of the net spectra through each of these bandpass filters, measured at the output of a single-mode fiber functionally identical to that shown in Fig. 3 (before crossing the vacuum feedthrough), are shown in Fig. 4. These spectra were measured using an Ocean Optics HR4000CG-UV-NIR spectrometer and calibrated relative to a 3200 K blackbody reference measured concurrently. The net spectrum through the 9.4% transmission bandpass filter has a similar spectral slope to a 1567 K blackbody (in an interval from 770 to 830 nm), with an arbitrary normalization factor (*i.e.*, similar in spectral slope but not brightness temperature). The FWHM of the net spectrum (measured f_λ) through this filter is 9.0%, smaller than the filter transmission FWHM because of the spectral slope. The effect of this spectral slope on performance is analyzed in Section 5.4.

The output of the SM fiber shown in Fig. 3, inside the vacuum chamber, is imaged with a 1:1 lens relay onto a 5 μm pinhole. This pinhole acts as a spatial filter, making the illumination in

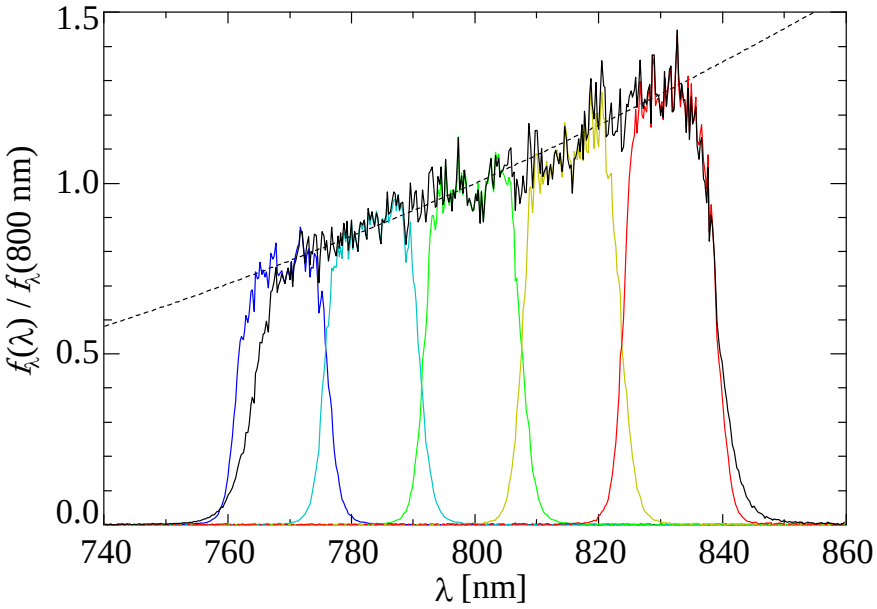

Fig. 4. Measured net spectra of supercontinuum source, through each of six bandpass filters (five 2% bandpasses and one 9.4%), with arbitrary normalizations. The black line is the 10% filter. The dotted line is a best-fit blackbody (with arbitrary normalization), with $T = 1567$ K. Each spectrum is a measured flux density (*i.e.*, units of [W/nm]), calibrated concurrently with an external 3200 K blackbody reference. These spectra are measured emerging from a single-mode fiber before passing the vacuum feedthrough into the tank.

the coronagraph pupil more uniform, and ensuring that the field surrounding the source is dark. The measured illumination of the coronagraph pupil shows a low-spatial-frequency intensity variation that is essentially quadratic in pupil radius, with the edges of the 30 mm pupil approximately 2% fainter than the center. The measured pupil illumination also shows high-spatial-frequency intensity variations of approximately 10% due to Fresnel diffraction of surface features on optics not located in pupil planes (*e.g.*, OAPs 1-4).

3.3 Elements of the Lyot Coronagraph

HCIT coronagraph configuration for this milestone demonstration is a Lyot coronagraph with a modified band-limited occulter. The concept of a band-limited coronagraph was introduced by Kuchner and Traub (2002). For Milestone 2, the occulter used is similar to a one-dimensional $(1-\mathrm{sinc}^2)$ fourth-order occulter, whose optical density (OD) profile (at 800 nm wavelength) is truncated and smoothed by convolution with a Gaussian function. This smoothed pattern is described in detail in Moody & Trauger (2007). Specifically, the sinc^2 intensity transmission profile is $T_{\mathrm{sinc}}(x) = \{1-[\sin(\pi x/w)/(\pi x/w)]^2\}^2$, $\mathrm{OD}_{\mathrm{sinc}}(x) = -\log_{10} T_{\mathrm{sinc}}(x)$, with $w = 142$ μm. The truncation and smoothing gives $\mathrm{OD}_{\mathrm{rel}}(x) = \min[\mathrm{OD}_{\mathrm{sinc}}(x), 8] \otimes G(x)$ with $G(x) = (2\pi\sigma^2)^{-1/2} \exp[-x^2/(2\sigma^2)]$, $\sigma = 9$ μm. For practical reasons, the maximum transmission is often less than unity, so the final transmission is $T(x) = T_0\ 10^{-\mathrm{ODrel}(x)}$, for some maximum transmission T_0. This transmission profile is shown graphically in Fig. 5.

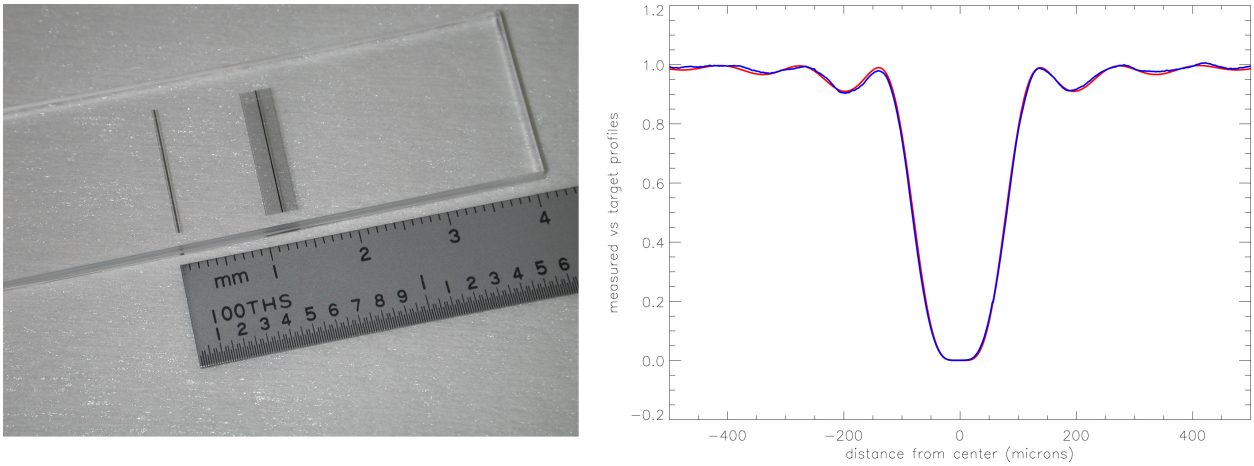

Fig. 5. (LEFT) Photo of Ni mask on fused silica substrate, (RIGHT) transmission profile of Ni mask as measured through microscope. In the transmission profile, the red line is the design profile, the blue line is the measured profile. Both transmission profiles have been normalized to unity at the peak transmission, *i.e.*, $T(x)/T_0$ is plotted.

For Milestone 2, the spatially varying transmission profile is optically realized by spatially varying thicknesses of Ni, deposited on a fused silica substrate. This occulter was fabricated at JPL, using electron-beam evaporation of Ni, with the substrate covered everywhere except for a moving 20 μm × 15 mm rectangular slit. The motion of the slit during fabrication is controlled at 1 μm increments, with μs dwell time precision. The measured transmission profile of the occulter used here is shown in Fig. 5.

Because Ni has a large index of refraction ($n \sim 2.5$ at 800 nm), regions of the occulter with higher OD (lower transmission, thicker Ni) also have a greater optical path length in transmission than low OD regions. The electric-field amplitude transmission is therefore complex-valued (with spatially varying modulus and phase), which changes the coronagraph behavior relative to a hypothetical occulter with real-valued transmission. In addition to the spatial variations in OD and phase, the OD and phase also vary with wavelength. Ni was chosen for this application because its OD and phase vary less with wavelength than other practical materials, as described in Balasubramanian (2008). It is worth noting that for two reasons, the mask described here is not a band-limited mask as described in Kuchner & Traub (2002); because by design its transmission is not zero for $x = 0$, and because its Fourier Transform does not have finite support (both because of the convolution with the Gaussian, and because the transmitted phase alters the real and imaginary amplitude transmission components).

The fused silica occulter substrate is 2.25 mm thick, with a multilayer dielectric antireflection coating on the face opposite the Ni profile. When installed in the coronagraph, it is tilted 5 degrees about the horizontal axis so that light reflected from the occulter can be collected by a beam-dump (seen in Fig. 2 as a reflected optical path upstream of the occulter).

The Lyot stop that accompanies this occulter is described by the intersection of two circles of diameter $D = 30$ mm, with centers separated by $\varepsilon D = 10.8$ mm (*i.e.*, $\varepsilon = 0.36$), shown in Fig. 6. This choice of ε does not arise from an exact correspondence with the width w of the occulter, as

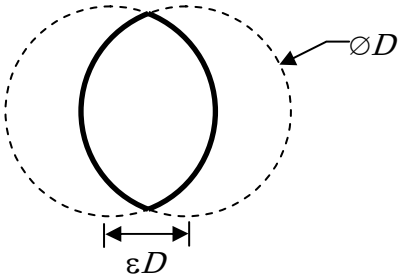

Fig. 6. Shape of Lyot stop. The clear aperture of the Lyot stop is the intersection of two circles of the same diameter as the pupil stop, D, with centers separated by εD (shown for $\varepsilon = 0.36$). Note that the εD measurement in the figure is between centers of the dashed circles, and is generally not equal to the linear width of the open aperture (*i.e.*, εD does not equal the width of the open aperture between the solid lines for $\varepsilon \neq 0.5$).

may be the case in a band-limited Lyot mask. This number instead is determined by simulations of the system and the wavefront correction, as described in Moody & Trauger (2007).

3.4 Wavefront sensing and control

While the speckle nulling wavefront control algorithm, as described in the M1 report and the M2 White Paper, was used successfully to demonstrate M1, a new algorithm has been implemented that reduces the time to create a dark hole and achieves better broadband performance. As such, the description of the speckle nulling algorithm will not be repeated here, as it was not used for M2. The algorithm used for wavefront control in this demonstration is the electric field conjugation (EFC) method, described in more detail in Give'on *et al.* (2007). The basic aspects are described here.

The EFC algorithm requires as an input an estimate of the complex *E*-field amplitude at the science camera, in the region over which corrections are desired, while a single intensity measurement at the science camera measures only the squared modulus of the *E*-field amplitude. The phase of the *E*-field must be determined before wavefront correction can be effective. Wavefront sensing for this algorithm is obtained in a manner analogous to phase-shifting interferometry (PSI), see *e.g.* Malacara (2005). The principle here is that for small changes in the DM surface height, the first-order effect of the DM is to add an electric field component in the pupil plane, and therefore to add some electric field component to the science camera image plane as well (as propagated through the coronagraph). From an initial DM setting, a series of DM changes are chosen to add, at each pixel in the sensing region of the image plane, a series of four (or any number not less than three) "probe" E-field components whose amplitudes are constant but whose phases change through the series. When each "probe" DM setting is applied to the DM, a science camera image is saved. The set of five images (one unprobed, four probed images) then determines, through standard PSI calculations, the complex value of the unprobed *E*-field at each pixel in the sensing region. This step is referred to as wavefront sensing, the estimation of complex image-plane *E*-fields.

Given this complex *E*-field estimate in the sensing region of the science camera image, and a model of the coronagraph as a whole (including the DM, occulter, Lyot stop, and optics), the

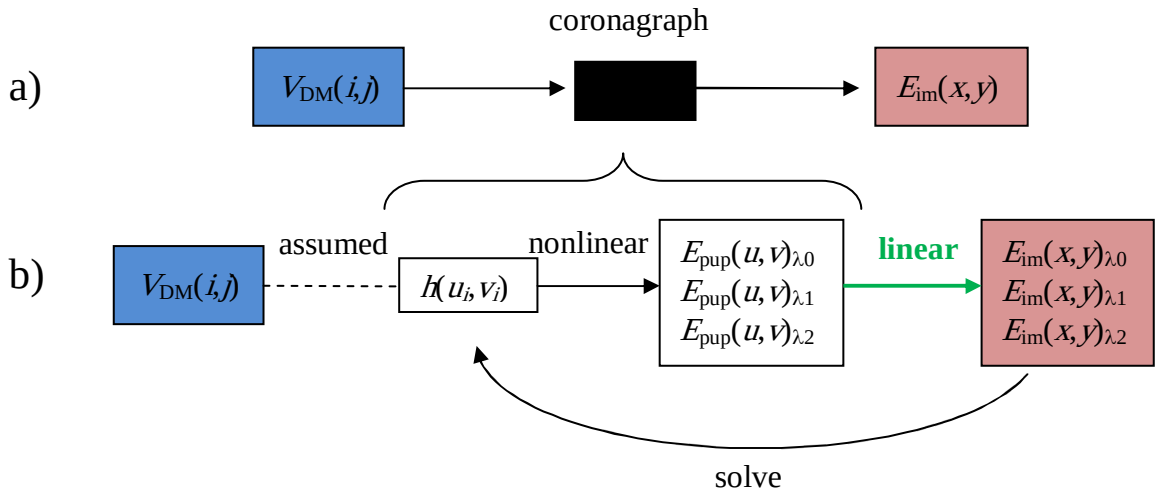

Fig. 7. Block diagram of EFC wavefront control algorithm. *A priori* model of coronagraph, including DM, occulter, Lyot stop, and optics defines a simple "black-box" model of DM control in 7a, taking 1024 DM voltages, $V_{DM}(i,j)$., and producing complex *E*-fields at each of ~ 4000 pixels in the image plane, $E_{im}(x,y)$. To simplify the calculations, we assume a relationship between DM voltages and actuator heights, $h(u_i,v_i)$, based on calibrations. This gives a straightforward influence on the complex $E_{pup}(u,v)$, but it must be calculated at each of a number of wavelengths, shown in 7b as λ_0, λ_1 and λ_2. The image plane representation is then a linear function of E_{pup}. EFC finds a nonlinear solution for *h* solving $|E_{im}| = 0$, which is overconstrained and yields a solution that minimizes the squared residuals, *i.e.*, the intensity.

wavefront control process becomes an exercise in nonlinear least squares optimization. Specifically, we have a system with DM voltage settings as inputs and science camera intensities as outputs, and we want to find the DM voltage settings that minimize the science camera intensities over some control region of the image. This is conceptually shown in Fig. 7.

For broadband performance at the level of M2, inputs from multiple wavelengths (*i.e.*, multiple bandpasses) are required. The mathematical formulation of this problem is then that we want to find a set of actuator heights, $h(u_i,v_i)$ that produces a zero *E*-field everywhere (both real and imaginary parts) in the controlled region at each wavelength. In practical terms, this is an attempt to use ~ 800 degrees of freedom (voltages on actuators that are not covered by the 30 mm pupil stop) to control ~ 12,000 image plane *E*-field values (2,000 CCD pixels, each with a real and an imaginary *E*-field component at each of 3 wavelengths). An exact solution will not be possible, so this is an optimization problem, finding the minimum squared residuals. The squared residuals are the science camera image intensities, exactly what we want to minimize. The unconstrained nonlinear least squares solution is a well-understood math problem, and we use an approach that is essentially a Levenberg-Marquardt algorithm, see *e.g.* Nocedal & Wright (2006). This is an iterative approach, which is well suited to application on a testbed where the real equipment differs at some level from the model used to determine the solution, and on which the DM may not act exactly as expected.

The fundamental flow of the wavefront sensing and correction is iterative. Each iteration proceeds as follows:

1) Take an image with the 10% bandpass filter to determine "current" contrast values.
2) For each of three 2% bandpass filters, centered at 768, 800, and 832 nm:
    a. Take an image at the "nominal" DM setting for this iteration.
    b. Take four "probe" images for four "probe" DM settings.
    c. Estimate complex $E_{im}(x,y)$ in the sensing region from 5 images for each λ.
3) Calculate and apply a new DM setting based on the three $E_{im}(x,y)$ maps.

In this scheme, no wavefront sensing is done using the 10% filter. The decision to use only three 2% filters for wavefront sensing and control instead of all five available filters was made to reduce the time to complete an iteration and seems to achieve a similar contrast in the end.

In practice, the total source throughput through each 2% bandpass filter is smaller than that of the 10% filter, and the exposure times are longer. In all, a single iteration including the 10% contrast measurement and the wavefront sensing in all three 2% filters takes a bit less than 40 minutes. This is significantly longer than the typical speckle nulling iterations described in the M2 White Paper (~ 7.5 minutes), but this algorithm requires far fewer iterations to reach a given contrast level than speckle nulling and reaches better ultimate contrasts in broadband light than speckle nulling. The relevance of this timing difference in the context of milestone performance will be discussed in Section 5, where it is shown that the milestone is achieved regardless of whether the requirements are interpreted as an elapsed time or as an accumulation of iterations.

4 MILESTONE PROCEDURES

Here we collect the various definitions, procedures, and requirements that comprise the Milestone 2 demonstration, as specified in the Milestone #2 White Paper (2006 Nov 28). The definitions of terms, in Section 4.1, are copied verbatim from the White Paper, except in the case of "speckle nulling," which has changed substantially, and is explained in its own section. The photometric procedures also differ somewhat from those described in the White Paper, and are described in their current form in Section 4.2, "Normalization of intensities to star brightness." Editorial comments to the verbatim text from the M2 White Paper are italicized.

4.1 Definition of Terms

The *M2* contrast metric requires a measurement of the intensity of speckles within the dark field relative to the intensity of the central star. Here we define the terms involved in the demonstration of M2.

Standard techniques for the acquisition of CCD images are used. We define a "raw" image to be the pixel-by-pixel image obtained by reading the charge from each pixel of the CCD, amplifying it, and sending it to an analog-to-digital converter. We define a "calibrated" image to be a raw image that has had background bias subtracted and the detector responsivity normalized by dividing by a flat-field image. Saturated images are avoided in order to avoid the confusion of CCD blooming and other potential CCD nonlinearities. All raw images are permanently archived and available for later analysis.

We define "scratch" to be a resting period of at least 30 minutes, throughout which time the DM actuators are held at equal voltages (e.g., 15 V on each actuator). The purpose of setting the DM to scratch is (a) to provide a neutral as well as standard starting point for subsequent DM operations, (b) to allow any transient physical effects to relax to a negligible level, and (c) to provide a clean separation boundary between previous DM settings and those to follow.

We define the "star" to be a small pinhole illuminated with broadband light relayed via optical fiber from a source outside the HCIT vacuum wall (*e.g.*, the super-continuum white light source). The "small" pinhole is to be unresolved by the optical system; e.g., a 5 μm diameter pinhole would be "small" and unresolved by the 21 μm FWHM Airy disk in an *F*/25 beam at 800 nm wavelength. This "star" is the only source of light in the optical path of HCIT. It is a stand-in for the star image that would have been formed by a telescope system in a space telescope.

We define the "algorithm" to be the computer code that takes as input the measured speckle field image, and produces as output a voltage value to be applied to each element of the DM, with the goal of reducing the intensity of speckles.

"Speckle nulling" had a particular definition at the time the White Paper was written, with the provision that HCIT may accommodate "New nulling algorithms for improved convergence rates and better performance in spectrally broad illumination." *The algorithm used for M2 is substantially different than the speckle nulling algorithm used for M1, and is described in detail in Section 3.4. The previous definition is not repeated here.*

"Statistical Confidence." The interpretation of measured numerical contrast values shall take into consideration, in an appropriate way, the statistics of measurement, including, for example, detector read noise, photon counting noise, and dark-count noise.

The goal is to demonstrate with high confidence that the true value of the contrast in the dark field, as estimated from our measurements (see Section 4.3), is less than the required threshold contrast value C_0. The estimated true contrast shall be obtained from the average of the set of approximately eight contrast values, each obtained with a new DM setting, measured in a one-hour period. For this milestone the required threshold is a mean contrast of $C_0 = 1.0 \times 10^{-9}$ with a confidence coefficient of 0.90 or better.

At any time in the demonstration, the HCIT contrast is subject to laboratory conditions, including the quality of the optical components, their alignment, any drift in their alignment over time, and the effectiveness of the speckle nulling wavefront sensing and control cycles. With each iteration our speckle nulling procedure attempts to improve the contrast metric, thus compensating for any drift or changes in alignment that may have occurred since the previous iteration, and variations may be expected due to experimental noise and any limitations in the algorithm. The distribution of contrast metrics at each iteration is regarded as random (Gaussian) about a mean contrast for the data set. In the unlikely event (no more than one in a given hour) of a failed iteration, that data outlier will be eliminated from the data set. We therefore consider the mean contrast as representative of the achieved contrast for a data set, and the distribution of contrast determinations among the eight DM settings for each set as a combination of both (random) speckle nulling variations and random measurement errors.

The mean contrast and confidence limits are computed in the following manner. The average of one or more images taken at each DM setting is used to compute the contrast metric. The mean contrast for a set of approximately eight DM settings taken in a given hour is:

$$\hat{c} = \sum_{i=1}^{n} \frac{c_i}{n}$$

where $n \approx 8$ is the number of individual DM settings and images in each one hour set. The standard deviation σ_{each} in the contrast values c_i obtained for individual DM settings within the set, which now includes both the measurement noise and the (assumed random) contrast variations due to changes in the DM settings for each speckle nulling iteration, is:

$$\sigma_{each} = \sqrt{\sum_{i=1}^{n} \frac{(c_i - \hat{c})^2}{n-1}}$$

Our estimate $\hat{c}$ is subject to uncertainty in the contrast measurements $\sigma_{mean} = \sigma_{each} / \sqrt{n}$ and the independently-determined overall errors in photometry σ_{phot}. With the assumption that the contrast values have a Gaussian distribution about the mean contrast, the statistical confidence that the mean contrast is less than $C_0 = 1 \times 10^{-9}$ is given by:

$$conf = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} e^{-z^2/2} dz$$

where $t = (C_0 - \hat{c}) / \sigma$ and $\sigma = \sqrt{\sigma_{mean}^2 + \sigma_{phot}^2}$. The values $\hat{c}$ and σ are the milestone metrics.

4.2 Normalization of intensities to star brightness

The absolute brightness levels measured at the science camera have no intrinsic significance, but the relative brightnesses of different parts of the image plane are significant, and most importantly, the brightnesses in the image plane with the coronagraph fully assembled should be referenced to the brightnesses in the image plane with the occulter laterally offset, as described below. The goal of interpreting measured brightnesses with an assembled coronagraph is then to normalize them to the peak brightness of the "unocculted star."

When imaging the source star through the coronagraph, with reference to Fig. 2, a "first image" of the source is formed at the occulter plane. The occulter need not be centered on this image. In a perfect system (perfect optical surfaces, no misalignments) with a band-limited Lyot coronagraph, the appearance of the image at the science camera (after propagating through the entire coronagraph) will have a point-spread function (PSF) defined by the shape of the Lyot stop, with total throughput determined by the intensity transmission of the occulter at the location where the first image lands. In this idealized (aberration-free, band-limited) condition, the PSFs for different star first image locations change only by total throughput, not by morphology. When the first image of the star lands where the occulter intensity transmission reaches its maximum value, T_0 (see Section 3.3), the science camera measures an image whose peak intensity defines the normalization factor to be applied to all subsequent images, regardless of location. The science camera image in this configuration is referred to as the "unocculted star" image. An example is shown in Fig. 8.

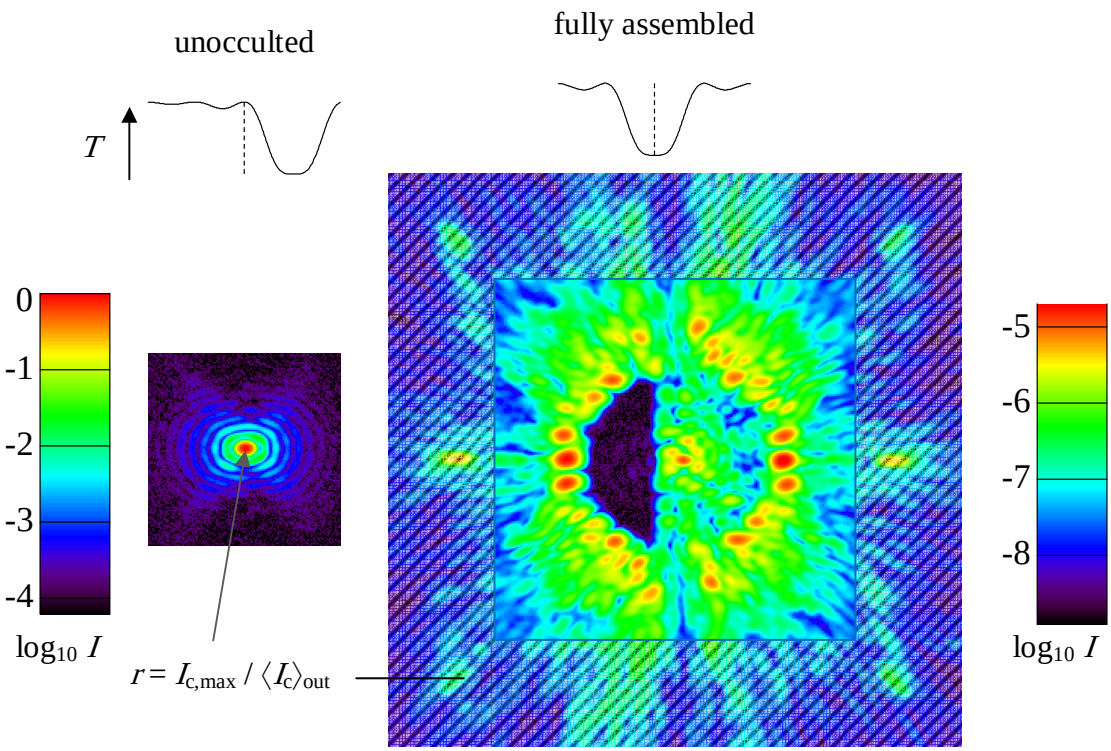

Fig. 8. Sample (LEFT) unocculted star and (RIGHT) assembled coronagraph science camera images. The hatched region on the right is the outer PSF region. The left colorbar shows normalized intensities of the unocculted image, the right colorbar shows normalized intensities of the fully assembled image. For a modest range of DM settings, the ratio of unocculted peak intensity to the outer PSF intensity with an assembled coronagraph is constant to ~ 3%. The curves above the images are occulter transmission plots, with a dotted line at the location of the first image.

In a real coronagraph, optical aberrations and deviations of the occulter from a band-limited profile change the measured PSFs. For example, with the first image of the star centered on the occulter, intensities in some parts of the science camera image reach 10^{-5} of the peak unocculted star brightness. This level of aberrations then defines the range of brightnesses that must be accommodated to compare unocculted images to assembled coronagraph images.

The supercontinuum source brightness varies for several reasons, notably the room temperature and the bandpass filter selection. It is then desirable to have a concurrent photometric reference in each science camera image, traceable to the peak measured intensity of the unocculted star image. The intensities of the outer regions of the PSF of the fully assembled coronagraph have proven empirically to vary with the full range of DM settings and other reasonable operating conditions (location of the first image of the star relative to the occulter, ambient temperature, *etc.*) only at the 10% level, and for the range of DM settings of interest during quasi-steady-state operation, to vary only at the 3% level. The goal of normalization is then to establish the ratio of peak unocculted star brightness to the brightness in a region of the outer PSF with the coronagraph fully assembled (*i.e.*, the first image of the star centered on the occulter), which is assumed to be constant to ~ 3%. This outer PSF region appears in every subsequent science

image (just as in Fig. 8), and provides exposure-by-exposure normalization to the unocculted star peak regardless of source brightness.

As mentioned above, the science camera image brightness with an assembled coronagraph reaches ~ 10^{-5} of the peak unocculted star brightness, requiring this level of dynamic range between measurements of unocculted and assembled images to establish the normalization. This dynamic range is too large to be accommodated by changing exposure times; assembled coronagraph images begin to saturate the CCD after a few seconds, so unocculted images would saturate in ~ 10 μs, too short for a mechanical shutter to accommodate.

The aperture stop in the supercontinuum source can stably attenuate the source output by as much as 10^{-9}, a larger factor than the 10^{-5} required for normalization. The net effect of the aperture stop can be represented as a reduction of source brightness $S = A\, S_0$, where S is the output source brightness, A is the attenuation from the aperture stop ($10^{-9} < A < 1$), and S_0 is the maximum source brightness. In addition, the location of the first image of the star on the occulter can be positioned to see any occulter intensity transmission T within the 10^{-5} practical range between fully assembled (with aberrations) and unocculted.

The practice adopted by HCIT to establish the intensity normalization is to begin with an aperture stop attenuation of $A_0 \sim 10^{-5}$, unocculted ($T = T_0$), with an exposure time chosen to avoid CCD saturation. The occulter is then moved so that the star location gives $T(x) = T_1 \sim 0.1\ T_0$ (10× less light at the camera), an image is taken, then the aperture stop is opened to $A_1 \sim 10^{-4}$ (*i.e.*, 10× brighter than A_0) and another image is taken. The exposure time is the same for all images, producing intensities that are either close to saturation or close to 0.1 × saturation. This cycle of decreasing T and increasing A is repeated until $T = T_{\min}$ and $A = 1$ (an unattenuated, fully assembled coronagraph). Fig. 9 gives a graphical representation of this sequence of images, requiring only two cycles to reach $T = T_{\min}$, $A = 1$, instead of the five cycles used in practice on HCIT.

For each pair of images at the same T_i (taken with A_{i-1} and A_i), a region of the image is chosen, in which sufficient light is present to use for calibration. The first image region, taken through T_0, is the single peak pixel. Given these regions, a sequence of ratios is calculated, relating the light from the image through T_i to the light in the image taken through T_{i+1}, both with A_i. A final ratio, taken with $T_{\min}$ and $A = 1$, references the "outer PSF" region (see Fig. 9). The ratios are defined as $r_i = \langle I_{\mathrm{c,norm}}(T_i, A_i)\rangle_i \,/\, \langle I_{\mathrm{c,norm}}(T_{i+1}, A_i)\rangle_{i+1}$, where $\langle\ldots\rangle_i$ is a mean over region i. The cumulative ratio $r = \Pi_{i=0}^{n} r_i$ is the normalization constant, $r = I_{\mathrm{c,norm,max}} \,/\, \langle I_{\mathrm{c,norm}}\rangle_{\mathrm{out}}$. In every subsequent assembled coronagraph image, the normalization is

$$I(x,y) = I_{\mathrm{c}}(x,y) \,/\, (r\,\langle I_{\mathrm{c}}\rangle_{\mathrm{out}}).$$

This normalization process is meaningful assuming that the source brightness does not significantly vary between two images taken with T_i and T_{i+1} but the same A_i. Any source variations between these images will produce an error in the measurement of r_i. In practice, multiple images are taken in each configuration, both to estimate and to mitigate this error. The source brightness fluctuations, while not carefully quantified, seem to exhibit the most power at low temporal frequencies (timescales longer than hours) rather than in the ~ 10 s timescales over

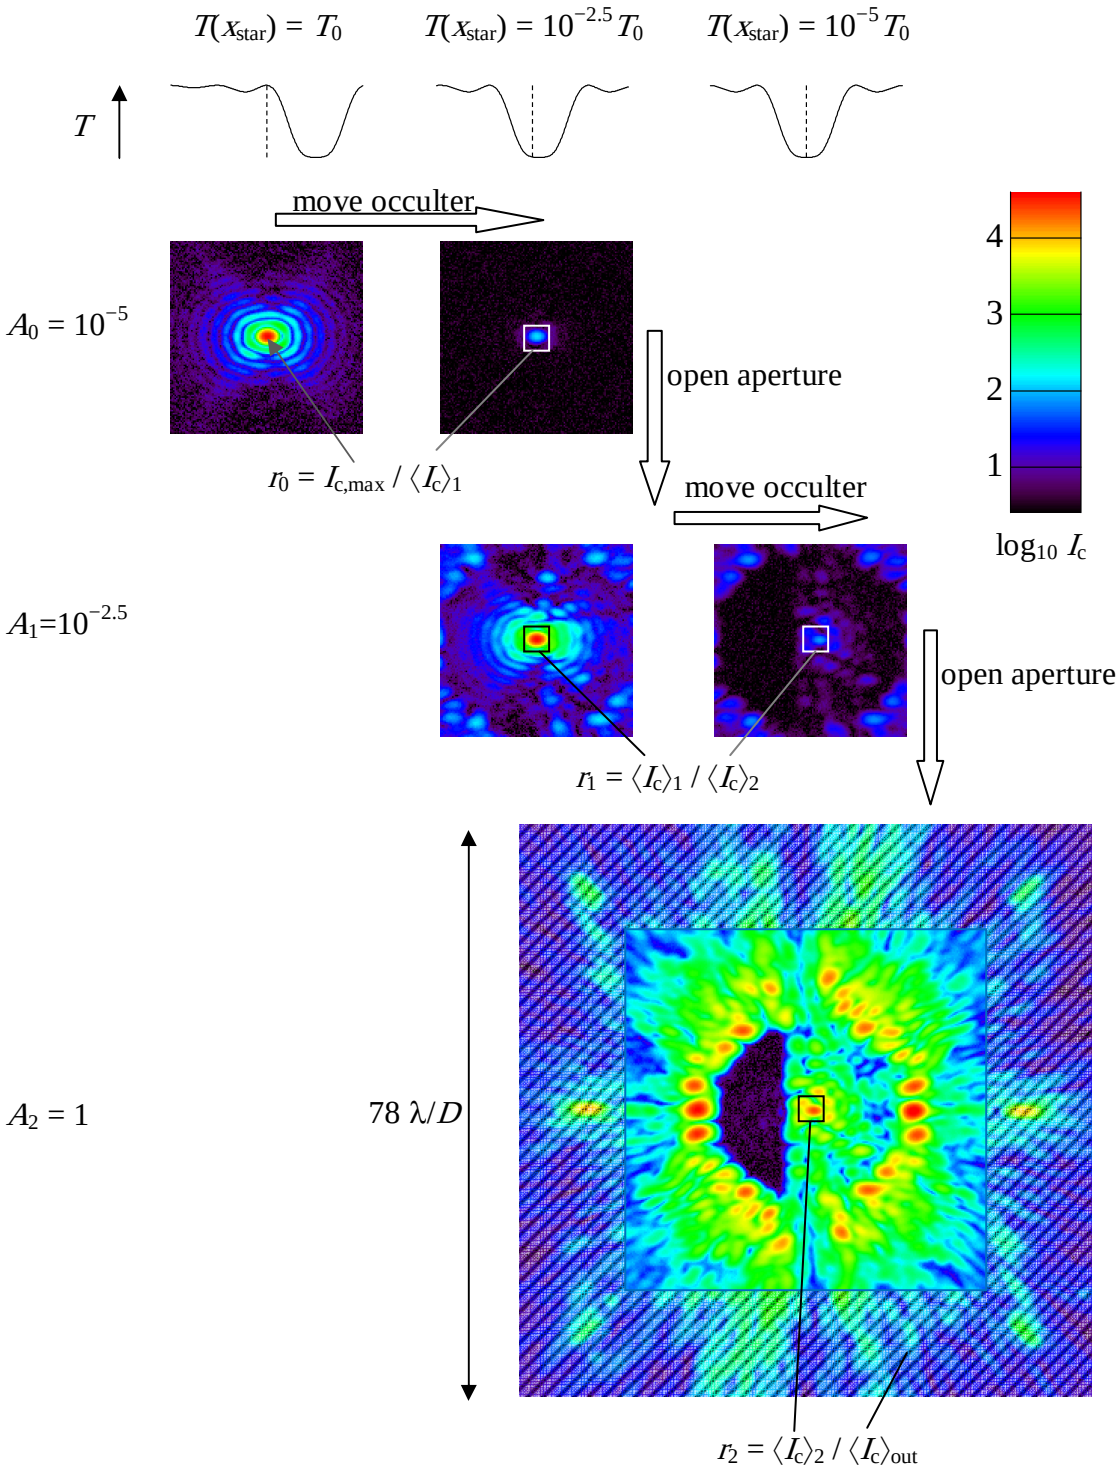

Fig. 9. Graphical representation of photometric calibration for image normalization, done in two stages. The rows reflect different source attenuations, A. The columns reflect the occulter position, *i.e.*, different offsets between the star and occulter, giving different levels of transmission, T. The regions of the images (black or white squares) in which intensities are calculated for the flux ratios are the same for images in the same column. The intensity scaling for all these images is shown in the color bar, measuring I_c, representing "detector numbers" (DN) at the science camera (*i.e.*, before normalization).

which this technique is sensitive to fluctuations. At a single DM setting, the entire calibration process is repeatable at the ~ 2% level.

The interpretation of normalized intensities is then that an image at the science camera has been normalized by a single scale factor so that if the occulter were instead "removed," with no other changes to the system, the peak intensity would have been unity. The full science camera PSF as measured is intrinsically related to the input star image as seen without an occulter (but with the Lyot stop in place). This normalization uses purely empirical measurements, which unambiguously establish the reference for all intensities, regardless of the actual shape of the PSF, the relationship between the aperture stop and the output brightness, the actual transmission profile of the occulter, or any nonlinearities in received flux as a function of exposure time.

4.3 Contrast Measurements

The normalized intensities obtained through the procedure described in Section 4.2 are an excellent description of the PSF of the star as seen through the assembled coronagraph, but do not directly describe how bright the wings of the star PSF would be relative to the core of a nearby planet PSF. The occulter intensity transmission, $T(x)$, will attenuate the planet PSF, so when referencing intensities to the brightness of a planet, the occulter transmission must be divided out. There are then two useful representations of the star PSF as seen through the assembled coronagraph: the normalized intensity of the PSF (described above in Section 4.2) and the contrast, described here.

Contrast is defined at each pixel as the relative brightness of a planet centered at that pixel whose PSF peak intensity would equal the star PSF intensity at that pixel. In other words, at a pixel where the star PSF contrast is 10^{-8}, the addition of a planet 10^{-8} as bright as the star centered at that pixel would double the measured intensity in that pixel. Algebraically, the contrast c is defined as

$$c(x,y) = I(x,y)\ T_0/T(x,y).$$

As a practical matter, because the difference between the designed and measured occulter transmission is not significant, the T/T_0 used to calculate c is the designed transmission profile as defined in Section 3.3.

The star PSF "contrast" as defined here has some relationship to a detectability threshold, depending on the observing strategy. For a single image, in which CCD read noise and shot noise were negligible, the star PSF contrast is approximately a unity signal-to-noise ratio (SNR) planet brightness threshold, which is an underestimate of the minimum brightness of a planet that could reasonably be detected. If multiple images can be taken with some offset (*i.e.*, a telescope roll angle), then the minimum detectable planet can be much fainter than the star PSF contrast, as demonstrated in Trauger & Traub (2007). Spectral information (*e.g.*, multiple bandpass filters) can similarly make fainter planets detectable in the presence of a finite star PSF contrast. This contrast metric should be interpreted as an input to a total error budget, not as a direct detectability threshold.

For the purposes of reducing the data to a manageable extent, for every iteration, two scoring regions are defined, and the pixel-by-pixel contrast values in each scoring region are averaged to

give two mean contrasts. The first scoring region is a $1 \times 1\ \lambda/D$ box, ranging from $4 < -x/(\lambda/D) < 5$ and $-0.5 < y/(\lambda/D) < 0.5$, named the "inner scoring region," corresponding to the inner working angle of interest. The second scoring region is defined by $-x > 4\ \lambda/D$, $r < 10\ \lambda/D$, a "D"-shaped region with area ~ 79 $(\lambda/D)^2$, called the "outer scoring region," representative of a large search area. The inner scoring region is a subset of the outer scoring region. In each iteration, all of the pixel-by-pixel contrast values $c(x,y)$ in the inner scoring region are averaged to give c_{in}, and $c(x,y)$ in the outer scoring region are averaged to give c_{out}.

4.4 Success Criteria

As in Section 4.1, editorial comments not present in the M2 White Paper are italicized.

1) Illumination is spectrally broadband with a bandwidth $\delta\lambda/\lambda_0$ of 10%, centered at a wavelength λ_0 in the range between 500 and 800 nm.
2) A mean contrast metric of 1×10^{-9} or smaller must be achieved in both an outer target dark area ranging from 4 to 10 λ_0/D and an inner area ranging from 4 to 5 λ_0/D.
3) Criteria 1 and 2, averaged over data obtained during a period of one hour or more, must be met with a high confidence of 90% or better. Typically a new contrast field image is obtained at the end of each sensing cycle, *(note original wording)* approximately every eight minutes, for a total of about 8 images per hour. *Since the time of the M2 White Paper, the algorithm has been modified and each iteration cycle takes 40 minutes. This is described in Sections 3.4 and 5.2.* The mean contrast obtained from the hour-long dataset constitutes our best estimate of the achieved contrast.
4) The above tests will be repeated to produce a total of 6 data sets, each starting from "scratch," with starting times distributed over a period of three days or more.

4.5 Certification Data Package

The following data package requirements are copied from the M2 White Paper, followed by a section identifier to locate these products in this report.

1) A narrative report, including a discussion of how each element of the milestone was met, an explanation of each image or group of images, appropriate tables and summary charts, and a narrative summary of the overall milestone achievement. *Section 5*
2) A description of the optical elements, their significant characteristics, and their layout in the HCIT. *Sections 3.1-3.3*
3) A tabulation of the significant operating parameters of the apparatus, including temperature stability. *Appendix 2*
4) A calibrated image of the reference star, and an estimate of photometry errors. *Fig. 8*
5) Calibrated images of the occulter transmittance pattern and/or the measured transmittance profile. *Section 3.3*
6) Spectrum of the broadband light and an estimate of the intensity uniformity and stability of the illumination reaching the defining pupil (at the DM). *Section 3.2*

7) A contrast field image representative, within error limits, of the super set of data, with appropriate numerical or color-coded or grey-scale coded contrast values indicated, and with coordinate scales indicated in units of Airy distance (λ_0/D). *Fig. 1*
8) A contrast metric value for each target area in each contrast field image, *i.e.*, the super set of data, in tabular form. *Appendix 1*
9) Statistical data on the super set, including the global average metric value and associated uncertainty, as well as a summary of the average scatter within each data set and among all images in super set. *Section 5.2, Tables 1-2*

5 NARRATIVE REPORT

Here we describe how each element of the milestone was met, with an explanation of the data obtained, analysis of the images and the contrast metric, illustrated with appropriate tables and summary charts, and a narrative summary of the overall milestone achievement. We provide graphs and tables of the contrast metric for the inner and outer target areas in each contrast field image.

5.1 Definition of Image Plane Regions

There are five regions of the image plane at the science camera, defined at various points in the preceding sections, that are of interest to this experiment. These regions are a probed region, a controlled region, an outer PSF for normalization, an inner scoring region, and an outer scoring region. These regions are illustrated in Fig. 10, except for the outer PSF region that appears in Figs. 8 and 9.

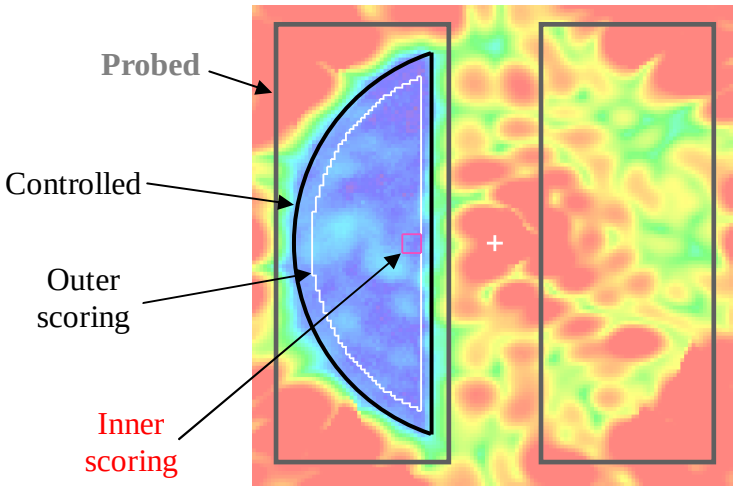

Fig. 10. Graphical description of image plane regions. The grey boxes outline the probed regions, where DM probes enable complex image-plane *E*-field estimation (wavefront sensing). The reverse-D shape outlined in black is the controlled region, where the wavefront control algorithm creates a dark hole. The reverse-D shape outlined in white is the outer scoring region, and the small box outlined in red is the inner scoring region, where contrast is compared to the milestone requirements.

The probed region is defined by the form of the DM probes applied during wavefront sensing (see Section 3.4), and was somewhat arbitrarily chosen to span $2.5 < |x/(\lambda/D)| < 12$, $-12 < y/(\lambda/D) < 12$. This region is essentially formed as a Fourier Transform of the differential DM surface change during probing, and represents that area over which the observed image plane intensities are modulated by the probes. As such, its "edges" in the image plane are smoothed by the finite number of actuators, and the image plane pattern (the region where the intensities change when probing) is symmetric by construction.

The controlled region is defined internally to the wavefront control algorithm, by choosing which image plane pixels to make corrections to. It too can have a "smoothed" edge, by applying different weights in the control algorithm to the corrections applied to individual pixels. A generic feature of the wavefront control algorithm used is that the "uncontrolled" pixels adjacent to the edges of the controlled region generally increase in intensity, so the controlled region is chosen to be a bit larger than the scoring regions to avoid scoring the "edge" effects.

The outer PSF is a region of the image that is relatively unaffected by changes in the DM settings. The presence of fine features (*i.e.*, high spatial frequencies) in the DM's shape changes determines the distance from the star to the outer PSF. One immediate implication is that the distance to the outer PSF must be greater than the corresponding Nyquist frequency, $15\ \lambda/D$. In practice, the location of the outer PSF used for normalization is chosen to be between 25 and $40\ \lambda/D$, which has empirically shown little variation (~ 3%) when changing DM settings.

The inner and outer scoring regions are defined in the M2 White Paper, repeated in Section 4.3. There is no smoothing of the edges of this region. Pixels are chosen to belong to the scoring region or not based on both the coordinate locations of the center of each pixel and the pixel scale of 4.87 pixels / (λ/D). The pixel where the unocculted star intensity peaks is defined to have a center at $(x,y) = (0,0)$, and every other pixel center differs in x and y by integer multiples of $(1/4.87)\ \lambda/D$. The inner scoring region contains 25 pixels, and the outer scoring region contains 1871 pixels. The average $c(x,y)$ contrast values over these pixels define c_{in} and c_{out}, the reduced contrast values for each iteration.

5.2 Data Acquisition

The requirements described in elements 3 and 4 of Section 4.4 set forth a timeline over which the contrast measurements were taken. Fig. 11 shows a graphical timeline of contrast measurements taken over 1 week, chosen somewhat arbitrarily from a longer data set of similar performance. In that week there were 7 acceptable contiguous data sets with good broadband contrast which were separated by DM scratch settings.

Table 1 lists the times during which the DM scratch settings were applied and the start and stop times of the M2 contrast data sets. Each scratch setting lasts more than 30 minutes, there are 7 data sets whose mean contrasts in the inner and outer scoring regions are below 10^{-9}, the duration of each data set is greater than 1 hour, the number of iterations in each data set is at least 8, and the total time between data sets is greater than 3 days.

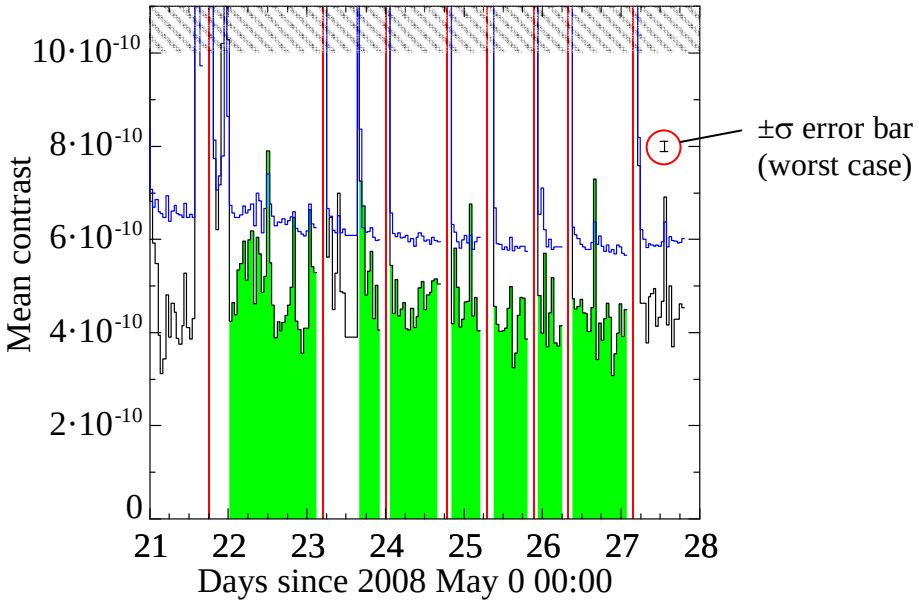

Fig. 11. Mean contrast per iteration, in the inner scoring region (black) and outer scoring region (blue), versus time. Red denotes time intervals when the DM was set to scratch. Green denotes intervals when at least 8 contiguous iterations each had both mean contrasts below 10^{-9}. The hatched region is $c > 10^{-9}$. The small error bar (circled) at the top right is the largest σ of all sets (see Table 2), dominated by calibration uncertainty.

By changing the wavefront sensing and control algorithm, the duration of a single iteration for M2 (~ 40 min) is very different than it was for M1 (~ 7 min). The White Paper timing requirements were then written to equate an hour of operation with eight iterations. Now that one hour of operation equates to 1.5 iterations, the timing context of the White Paper may be interpreted to place significance on one hour or on eight iterations. We show in the next section that the contrast requirements are met under either interpretation.

5.3 Image Analysis and Contrast Metrics

Section 4.1 defined the statistics to be examined when analyzing science images. In particular, estimates of σ_{mean} and σ_{phot} define the limits of a 90% confidence interval, considered here as a success criterion. We have estimated from historical data that the photometric calibration changes by less than 3%, so we adopt $\sigma_{phot} = 0.03\ \hat{c}$ as the rms error for a set, a conservative estimate. The estimates of σ_{mean} come from the individual data sets. The values tabulated in Table 2 conform to the interpretation of M2 requiring at least 8 iterations per set. In all cases, the photometric uncertainty dominates the total uncertainty. Under the assumptions laid out in Section 4.1 (specifically, Gaussian errors), a one-sided 90% confidence interval can be estimated around the measured mean contrast in each region. The upper limit to this 90% confidence interval, c_{90} (appearing in the last column of Table 2), is the final discriminator used by the M2 White Paper as a criterion for success. Our margin of success is measured by the degree to which the c_{90} values for inner and outer regions, over all data sets, are all well below 10^{-9}. When using the alternative interpretation on iterations required for M2, using only two iterations per set, the c_{90} numbers are lower than in Table 2 and are not tabulated separately. There were no "failed iterations" as defined in the "statistical confidence" description of Section 4.1.

	Start time (UTC)	Stop time (UTC)	Duration	Num iterations	$\hat{c}_{in} \times 10^{10}$	$\hat{c}_{out} \times 10^{10}$
Scratch 1	2008 May 21 17:50:15	2008 May 21 18:24:06	0:33:51			
Set 1	2008 May 22 01:06:48	2008 May 23 02:51:32	25:44:44	33	5.0	6.5
Scratch 2	2008 May 23 04:39:06	2008 May 23 05:10:55	0:31:49			
Set 2	2008 May 23 16:51:19	2008 May 23 22:24:22	5:33:03	8	5.4	6.4
Scratch 3	2008 May 23 23:55:09	2008 May 24 00:26:55	0:31:46			
Set 3	2008 May 24 02:09:59	2008 May 24 15:45:19	13:35:20	18	4.7	6.1
Scratch 4	2008 May 24 18:37:50	2008 May 24 19:14:20	0:36:30			
Set 4	2008 May 24 20:56:25	2008 May 25 04:52:50	7:56:25	11	4.8	6.0
Scratch 5	2008 May 25 06:42:15	2008 May 25 07:17:28	0:35:13			
Set 5	2008 May 25 09:50:01	2008 May 25 19:29:20	9:39:19	13	4.2	5.9
Scratch 6	2008 May 25 21:05:12	2008 May 25 21:42:05	0:36:53			
Set 6	2008 May 25 23:33:56	2008 May 26 05:54:14	6:20:18	9	4.4	6.1
Scratch 7	2008 May 26 07:33:19	2008 May 26 08:08:51	0:35:32			
Set 7	2008 May 26 09:47:31	2008 May 27 01:46:16	15:58:45	21	4.4	5.9

Table 1. Timing of DM resets and milestone data set acquisition.

Over the entire dataset, the global mean contrasts are 4.7×10^{-10} and 6.2×10^{-10} for the inner and outer regions.

Fig. 11 (and Table 2) shows a trend toward lower contrast over time. This is most clearly seen in the contrast values of the outer scoring region, which have less measurement noise. This has two immediate consequences: with more iterations, it is reasonable to expect better contrast numbers, and the σ_{mean} estimates are an overly conservative measure of the current performance because the variations from iteration to iteration in the measured data are not independent of one another. Both of these consequences support the conclusion that the results presented here are in no way marginal results and that 10^{-9} contrasts have been achieved.

The specificity of the statistical requirements in the M2 White Paper was an outgrowth of the need to define exactly the acceptability of the M1 performance. The M1 results measured

	$\hat{c}\times10^{10}$	$\sigma_{each}\times10^{10}$	$\sigma_{mean}\times10^{10}$	$\sigma_{phot}\times10^{10}$	$\sigma\times10^{10}$	$c_{90}\times10^{10}$
Set 1, inner	5.03	0.17	0.03	0.15	0.15	5.22
outer	6.50	0.05	0.01	0.20	0.20	6.75
Set 2, inner	5.40	0.40	0.14	0.16	0.21	5.68
outer	6.40	0.28	0.10	0.19	0.22	6.68
Set 3, inner	4.67	0.10	0.02	0.14	0.14	4.85
outer	6.06	0.03	0.01	0.18	0.18	6.29
Set 4, inner	4.79	0.25	0.07	0.14	0.16	5.00
outer	6.01	0.05	0.01	0.18	0.18	6.24
Set 5, inner	4.23	0.14	0.04	0.13	0.13	4.40
outer	5.91	0.07	0.02	0.18	0.18	6.14
Set 6, inner	4.38	0.24	0.08	0.13	0.15	4.58
outer	6.11	0.15	0.05	0.18	0.19	6.36
Set 7, inner	4.36	0.18	0.04	0.13	0.14	4.54
outer	5.90	0.04	0.01	0.18	0.18	6.12

Table 2. Statistics for data sets, each encompassing all iterations (not less than 8) listed in Table 1 (*i.e.*, coded in green in Fig. 11).

$c_{in} \sim 9\times10^{-10}$ with some iterations over 10^{-9}. The marginality of this data demanded an exact definition of the statistical measure of success. The M2 data satisfy the requirements with an appreciable margin (by nearly a factor of 2 in mean contrast). The statistical analysis presented here describes this adequately, but is somewhat superfluous given that Fig. 11 shows such a clear difference between the requirement and the measurements. The improvement from M1 to M2 is mostly attributable to the improvements in the wavefront control algorithm, which allows more rapid convergence and better final numbers than the speckle nulling algorithm used for M1.

5.4 Wavelength dependence

Images are taken in three of the 2% bandpass filters (see Fig. 4) to perform wavefront sensing. The complex *E*-field estimates are used by the correction algorithm to determine the appropriate wavefront control in each iteration. These three controlled bandpasses, along with the two uncontrolled 2% bandpasses (see Section 3.2), provide five spectral resolution elements covering the 10% bandpass defined by M2, from 760 to 840 nm. These images, through all six bandpass filters (one 9%, five 2% bandpasses) in one iteration, are shown in Fig. 12.

The images taken in each 2% filter are independently normalized as described in Section 4.2. A sum of these 2% images, with some weighting of each bandpass, represents the image of a star with a synthetic spectrum determined by the weights and the measured bandpass net spectra (from Fig. 4). The contrast in the inner and outer scoring regions is then the mean of the 2% contrasts, weighted by the same weights that define the synthetic star spectrum. Any chosen set of weights defines both a synthetic star spectrum and a synthetic contrast (integrated over the bandpass).

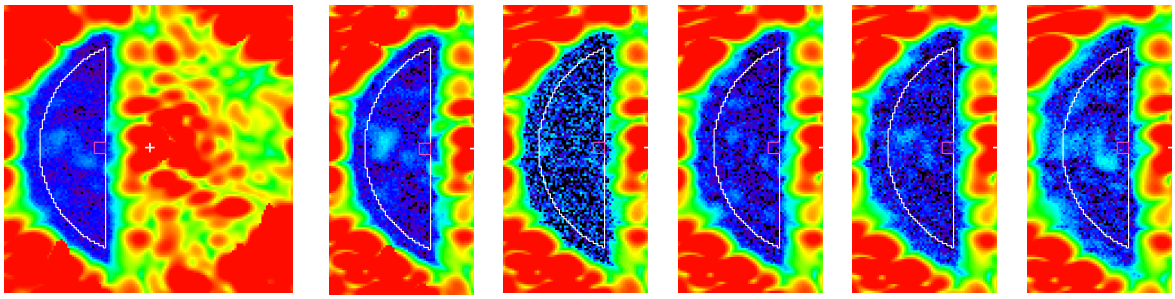

9%, 800 nm 2%, 768 nm 784 nm 800 nm 816 nm 832 nm

Fig. 12. Normalized intensity images across six bandpass filters. The bandpasses of the filters are listed below the images. This is iteration #12 from set #3.

Contrast as a function of wavelength, taken from the same iteration as Fig. 12, is shown in Fig. 13. The contrast in each 2% bandpass is well below 10^{-9}. It must then be true that any weighted mean of these contrasts with nonnegative weights will also be below 10^{-9}, which implies that any synthetic star spectrum with nonnegative weights will meet the M2 mean contrast requirement.

The first synthetic star spectrum to be analyzed quantitatively is formed by an equally weighted sum of the five independently normalized 2% bandpasses. This equal weight spectrum is shown in Fig. 14. The equal weight spectrum has a FWHM of 9.7%, broader than the 9.0% FWHM net spectrum from the broadband filter (Fig. 4). The contrast values for this equal weight spectrum, which are the mean of the measured 2% contrasts, over the entire data set are shown in Table 3. They are nearly equal to the contrasts measured through the 9% filter.

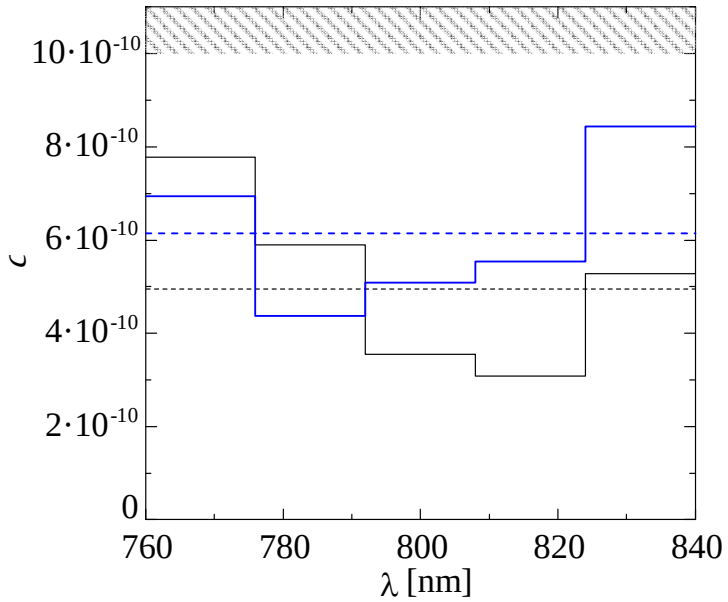

Fig. 13. Contrast vs. wavelength, same data as Fig. 12. The bandpasses of filters are represented here as "square" bandpasses. Black lines are c_{in}, blue lines are c_{out}. Solid lines are 2% bandpasses, dotted lines are the 9% bandpass. The hatched region is $c > 10^{-9}$.

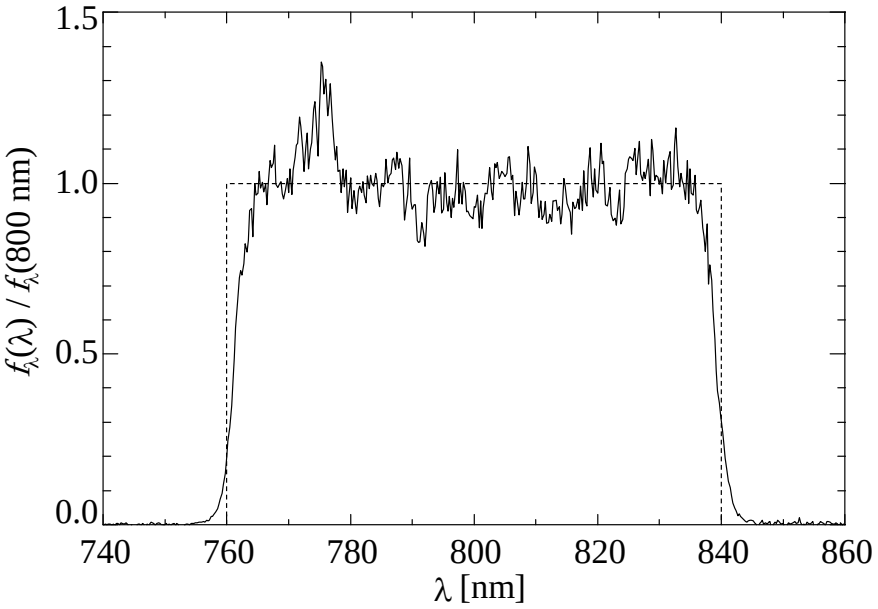

Fig. 14. Synthetic star spectrum formed by equal weights on five independently normalized 2% bandpasses (solid line) and flat spectrum (dotted line). The bump at 775 nm is where two neighboring filters overlap more than at other boundaries. The equal-weight FWHM is 77.5 nm, or 9.7%, the flat FWHM is 10.0%.

The second synthetic star spectrum to be analyzed quantitatively is a perfectly flat spectrum (in f_λ), with exactly a 10% fractional width, also shown in Fig. 14. An exact calculation of the contrast that would be measured from this spectrum would require knowledge of the contrast as a unction of wavelength with the same spectral resolution as Fig. 4 (approximately 0.25 nm). We estimate this relationship using a quadratic fit of contrast vs. wavelength (see Fig. 13), assuming that the measured 2% contrasts were measured from stars with exactly the net 2% spectra shown in Fig. 4. Using these quadratic fits (to five contrast measurements for each inner / outer region in each iteration), the contrasts are integrated over the flat spectrum to estimate the observable flat 10% contrast. These numbers are tabulated in Table 3, and are all within 10^{-11} of the equal weight contrasts.

	Equal $\hat{c}_{in}$	$\hat{c}_{out}$	$c_{in,90}$	$c_{out,90}$	Flat $\hat{c}_{in}$	$\hat{c}_{out}$	$c_{in,90}$	$c_{out,90}$
Set 1	5.5	6.6	5.7	6.8	5.5	6.6	5.7	6.9
Set 2	4.9	6.3	5.1	6.5	4.9	6.3	5.0	6.6
Set 3	4.7	6.0	4.8	6.3	4.7	6.1	4.8	6.3
Set 4	4.7	6.1	4.8	6.3	4.7	6.1	4.8	6.4
Set 5	4.0	5.9	4.2	6.1	4.0	5.9	4.2	6.1
Set 6	4.3	6.1	4.5	6.3	4.3	6.1	4.5	6.4
Set 7	4.5	5.8	4.6	6.1	4.5	5.9	4.6	6.1

Table 3. Inferred contrast values for synthetic star spectra, multipled by 10^{10}. "Equal" contrasts represent an equal weight synthetic spectrum (solid line in Fig. 14), "Flat" contrasts represent a flat synthetic spectrum (dotted line in Fig. 14).

6 CONCLUSIONS

This report described the equipment and procedures used to control coronagraphic images at levels of interest to flight missions, both in overall contrast level and in spectral bandwidth. Our results meet the requirements set out in the Milestone 2 White Paper to demonstrate this performance.

Table 2 contains the statistical analysis of the 9% fractional bandwidth contrast data, showing that the measurements presented here are better than 10^{-9} in the inner and outer regions of the images. Table 1 shows that the correction algorithm can maintain this contrast level over many hours, and even over many days if not continuing to reset the DM periodically. Table 3 shows that careful analysis of the spectra imply that the full 10% fractional bandwidth would produce contrasts well below the milestone requirements.

ACKNOWLEDGEMENTS

The work described in this report was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

APPENDIX 1: TABULATED CONTRAST VALUES

Table 4 contains the c_{in} and c_{out} values multiplied by 10^{10} for every iteration from Table 1.

Set	Iter.	c_{in}	c_{out}	Iter.	c_{in}	c_{out}	Iter.	c_{in}	c_{out}	Iter.	c_{in}	c_{out}
1	1	4.2	6.7	2	4.6	6.6	3	4.4	6.5	4	5.3	6.5
	5	5.5	6.5	6	6.0	6.7	7	5.1	6.6	8	6.0	6.6
	9	6.2	6.8	10	4.6	6.3	11	6.1	7.0	12	5.7	6.8
	13	4.9	6.1	14	5.2	6.7	15	7.9	7.4	16	5.5	6.8
	17	4.6	6.5	18	3.9	6.3	19	4.2	6.4	20	4.0	6.4
	21	4.2	6.5	22	4.4	6.3	23	4.6	6.4	24	5.0	6.5
	25	6.5	6.6	26	4.2	6.2	27	4.1	6.2	28	3.5	6.1
	29	4.1	6.1	30	4.1	6.2	31	6.6	6.8	32	5.4	6.3
	33	5.3	6.3									
2	1	7.3	8.4	2	6.7	6.3	3	4.8	6.2	4	5.3	6.2
	5	5.7	6.3	6	4.3	6.1	7	5.0	6.0	8	4.1	6.0
3	1	5.4	6.6	2	4.4	6.1	3	5.1	6.1	4	4.4	6.1
	5	4.5	6.0	6	4.6	6.1	7	4.1	6.1	8	4.1	6.0
	9	4.5	6.0	10	4.1	5.9	11	4.3	6.0	12	5.0	6.1
	13	5.1	6.1	14	4.5	6.0	15	4.8	6.0	16	4.9	5.9
	17	5.1	6.1	18	5.2	6.0						
4	1	4.2	6.3	2	5.8	6.2	3	5.0	6.0	4	4.1	5.8
	5	4.3	6.0	6	4.7	6.1	7	4.7	5.9	8	6.8	6.1
	9	4.4	5.8	10	4.8	5.9	11	4.0	6.0			
5	1	4.6	6.7	2	4.2	5.9	3	4.0	5.9	4	4.0	5.8
	5	4.1	5.8	6	4.5	5.8	7	5.0	6.1	8	3.2	5.8
	9	3.6	5.9	10	4.4	5.8	11	4.8	5.8	12	4.7	5.9
	13	3.9	5.7									
6	1	4.8	6.5	2	4.0	7.1	3	5.7	6.2	4	3.7	5.8
	5	4.4	6.0	6	5.2	5.8	7	3.8	5.8	8	3.7	5.8
	9	4.1	5.8									
7	1	4.7	6.3	2	4.5	6.1	3	4.6	6.0	4	4.7	5.9
	5	4.4	5.8	6	4.4	5.8	7	4.0	5.8	8	4.5	5.9
	9	7.3	6.4	10	3.4	5.9	11	4.2	6.1	12	3.8	5.8
	13	4.3	5.8	14	4.6	5.9	15	4.3	5.7	16	3.1	5.8
	17	3.5	5.7	18	4.0	5.9	19	4.6	5.8	20	3.9	5.7
	21	4.5	5.7									

APPENDIX 2: THERMAL ENVIRONMENT

HCIT contrast performance depends on some environmental conditions, most notably temperature. Approximately 40 temperature readings are logged continuously during normal operation of HCIT, from sensors on the tank, table, DM, optics, camera, and room. The degree of correlation of any performance metric to any of these sensor readings has not been quantified to date. Fig. 15 is a plot of the temperature readings on the table and on the DM, to allow some general characterization of the stability of the thermal environment.

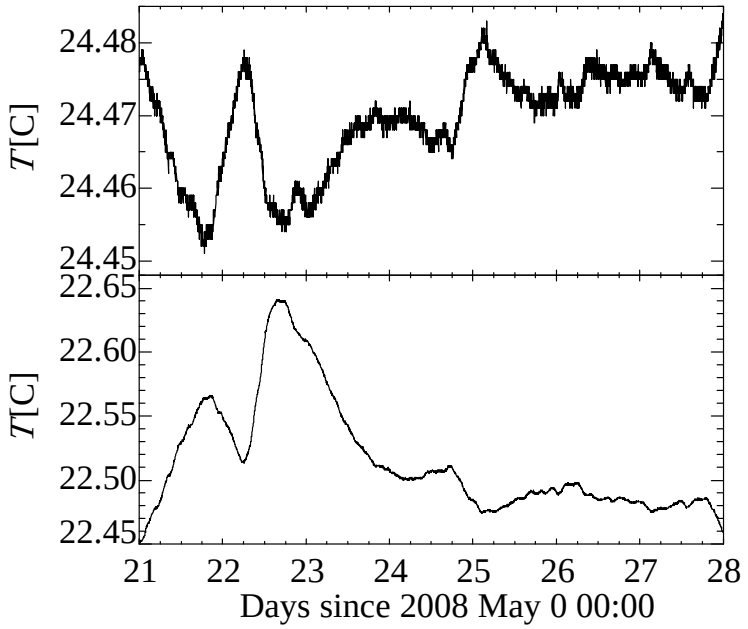

Fig. 15. Temperature readings from (TOP) DM and (BOTTOM) table top, over the same time period as Fig. 11. The T scales are different for the two plots. This table temperature sensor is located in the top-right corner of Fig. 2, where the coronagraph is most sensitive to motion. There is an obvious anticorrelation between these two temperatures. The rms fluctuations over the full week of readings are 7 mK for the DM and 44 mK for the table.

REFERENCES

Alfano, R.R. and Shapiro, S.L. 1970a, "Emission in the region 4000 to 7000 Å via four-photon coupling in glass," Phys. Rev. Lett. **24**, 584-587.

Alfano, R.R. and Shapiro, S.L. 1970b, "Observation of self-phase modulation and small-scale filaments in crystals and glasses," Phys. Rev. Lett. **24**, 592-594.

Balasubramanian, K. 2008, "Band-limited image plane masks for the Terrestrial Planet Finder coronagraph: materials and designs for broadband performance," Appl. Opt. **47**, 116-125.

Give'on, A., Kern, B., Shaklan, S., Moody, D.C. 2007, "Broadband wavefront correction algorithm for high-contrast imaging systems," Proc. SPIE 6691, 66910A1-11.

Kuchner, M.J. and Traub, W.A. 2002, "A coronagraph with a band-limited mask for finding terrestrial planets," ApJ **570**, 900-908.

Malacara, D., Servín, M., Malacara, Z., *Interferogram Analysis for Optical Testing: Second Edition*, New York:Taylor & Francis (2005).

Moody, D.C. and Trauger, J.T. 2007, "Hybrid Lyot coronagraph masks and wavefront control for improved spectral bandwidth and throughput," Proc. SPIE 6693, 6693I1-9.

Nocedal, J., Wright, S.J., *Numerical Optimization*, New York: Springer (2005).